



Attorney Docket: 3036/50061
PATENT

Applicant: SIMON PAUL DAVIS ET AL.

Serial No.: 09/898,484 Group Art Unit:

Filed: JULY 5, 2001 Examiner:

Title: IMPROVEMENTS IN OR RELATING TO SWITCHING DEVICES

CLAIM FOR PRIORITY UNDER 35 U.S.C. §119

Commissioner for Patents
Washington, D.C. 20231


Sir:

The benefit of the filing date of prior foreign application numbers 0016475.6 and 0024736.1, filed in Germany on July 5, 2000 and October 10, 2000 respectively, is hereby requested and the right of priority under 35 U.S.C. §119 is hereby claimed.

In support of this claim, filed herewith is a certified copy of the original foreign application.

Respectfully submitted,

October 22, 2001



Gary R. Edwards
Registration No. 31,824

CROWELL & MORING, LLP
P.O. Box 14300
Washington, DC 20044-4300
Telephone No.: (202) 624-2500
Facsimile No.: (202) 628-8844



THIS PAGE BLANK (USPTO)

Europe



INVESTOR IN PEOPLE



CERTIFIED COPY OF PRIORITY DOCUMENT

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.


Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

W. Evans

Dated 13 June 2001

THIS PAGE BLANK (USPTO)



1000700 3374374-1 000393
 P91/7790 0.00-0024736.1

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
 Newport
 Gwent NP9 1RH

1. Your reference 2000P04864/GB/R76/MM/rr

2. Patent application number **0024736.1**
 (The Patent Office will fill in this part) **10 OCT 2000**

3. Full name, address and postcode of the or of each applicant (underline all surnames) ROKE MANOR RESEARCH LIMITED
 OLD SALISBURY LANE
 ROMSEY, HAMPSHIRE, SO51 0ZN
 UNITED KINGDOM

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation UNITED KINGDOM
5618455007

4. Title of the invention IMPROVEMENTS IN OR RELATING TO SWITCHING DEVICES

5. Name of your agent (if you have one) DEREK ALLEN
 "Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode) Siemens Shared Services Limited
 Intellectual Property Department
 Siemens House, Oldbury
 Bracknell, Berkshire RG12 8FZ
 United Kingdom

Patents ADP number (if you know it)

7396419002

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority application number (if you know it)	Date of filing (day / month / year)
	GB	0016475.4	05/07/2000

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application	Date of filing (day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

a) any applicant named in part 3 is not an inventor, or

b) there is an inventor who is not named as an applicant, or

c) any named applicant is a corporate body.

See note (d)) YES

Patents Form 1/77

Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form	0
Description	13
Claim(s)	3
Abstract	1
Drawing(s)	3

10. If you are also filing any of the following, state how many against each item.

Priority documents	
Translations of priority documents	
Statement of inventorship and right to grant of a patent (Patents Form 7/77)	3
Request for preliminary examination and search (Patents Form 9/77)	1
Request for substantive examination (Patents Form 10/77)	
Any other documents (please specify)	

11. I/We request the grant of a patent on the basis of this application.

Signature

Margaret Mackett

Date 06.10.2000

12. Name and daytime telephone number of person to contact in the United Kingdom

Margaret Mackett	01344 396808
------------------	--------------

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

IMPROVEMENTS IN OR RELATING TO SWITCHING DEVICES

The present invention relates to improvements in or relating to switching devices, and is more particularly concerned with a system for providing distributed schedules for such a device. In particular, the invention relates to the distribution of control between a management card and line interface cards (LICs) on a switching device. For the purposes of the following description the term switching device refers to any device which performs the function of a circuit switch, packet switch or a router.

Data is transferred over the Internet by means of a plurality of routing devices in accordance with a standard protocol known as Internet Protocol (IP). IP is a protocol based on the transfer of data in variable sized portions known as packets. All network traffic involves the transportation of packets of data. Routers are devices for accepting incoming packets; temporarily storing each packet; and then forwarding the packets to another part of the network.

Traffic volume in the Internet is growing exponentially, almost doubling every 3 months, and the capacity of conventional IP routers is insufficient to meet this demand. There is thus an urgent requirement for products that can route IP traffic at extremely large aggregate bandwidths in the order of several terabits per second. Such routing devices are termed "terabit routers".

Terabit routers require a scalable high capacity communications path between the point at which packets arrive at the router (the "ingress") and the point at which the packets leave the router (the "egress").

The packets transferred in accordance with IP can (and do) vary in size. Within routers it has been found useful to pass data in fixed sized

units. In routers, the data packets are partitioned into small fixed sized units, known as cells.

One suitable technique for implementing a scalable communications path is a backplane device, known as a cell based cross-bar. Data packets
5 are partitioned into cells by a plurality of ingress means for passage across the cross-bar.

The plurality of ingress means provide respective interfaces between incoming communications channels carrying incoming data and the backplane device. Similarly a plurality of egress means provide respective
10 interfaces between the backplane device and outgoing communications channels carrying outgoing data.

A general terabit router architecture bears some similarity to conventional router architecture. Packets of data arrive at input port(s) of ingress means and are routed as cells across the cross-bar to a
15 predetermined egress means which reassembles the packets and transmits them across its output port(s). Each ingress means maintains a separate packet queue for each egress means.

The ingress and egress means may be implemented as line interface cards (LICs). Since one of the functions regularly undertaken by the
20 ingress and egress means is forwarding, LICs may also be known as 'forwarders'. Further functions include congestion control and maintenance of external interfaces, input ports and output ports.

In a conventional cell based cross-bar each ingress means is connected to one or more of the egress means. However, each ingress
25 means is only capable of connecting to one egress means at any one time. Likewise, each egress means is only capable of connecting to one ingress means at a time.

All ingress means transmit in parallel and independently across the cross-bar. Furthermore cell transmission is synchronised with a cell cycle, having a period of, for example, 108.8ns.

5 The ingress means simultaneously each transmit a new cell with each new cell cycle.

The pattern of transmissions from the ingress means across the cross-bar to the egress means changes at the end of every cell cycle.

The co-ordination of the transmission and reception of cells is performed by a cross-bar controller.

10 A cross-bar controller is provided for efficient allocation of the bandwidth across the cross-bar. It calculates the rates that each ingress means must transmit to each egress means. This is the same as the rate at which data must be transmitted from each packet queue. The calculation makes use of real-time information, including traffic measurements and
15 indications from the ingress means. The indications from the ingress means include monitoring the current rates, queue lengths and buffer full flags. The details of the calculation are discussed more rigorously in the copending UK Patent Application Number 9907313.2 (docket number F21558/98P4863).

20 The cross-bar controller performs a further task; it serves to schedule the transfer of data efficiently across the cross-bar whilst maintaining the calculated rates. At the end of each cell cycle, the cross-bar controller communicates with the ingress and egress means as follows. Firstly, the cross-bar controller calculates and transmits to each ingress means the
25 identity of the next packet queue from which to transmit. Secondly, the cross-bar controller calculates and transmits to each egress means the identity of the ingress from which it must receive.

The system described above does have a number of disadvantages however. The cross-bar controller is responsible for controlling the cell cycle-by-cell cycle behaviour of each ingress and egress means. At the rates required by a terabit router, this amounts to demanding complex
5 hardware to implement the cross-bar controller, the ingress and the egress means. Furthermore the demand for higher capacity places stringent delay performance conditions upon the communication channels between the ingress and egress means and the cross-bar controller means.

When developing a system for particular traffic conditions, it is
10 disadvantageous to have to replace inappropriate hardware.

It is therefore an object of the invention to obviate or at least mitigate the aforementioned problems.

In accordance with one aspect of the present invention, there is provided a switching arrangement having:-

- 15 a cross-bar;
- a plurality of ingress means connected to an input side of the cross-bar, each ingress means including an ingress schedule storing means;
- a plurality of egress means connected to an output side of the cross-bar, each egress means including an egress schedule storing means; and
- 20 a management card which communicates configuration primitives to each of the plurality of ingress means and to each of the plurality of egress means, the configuration primitives providing updated entries for ingress and egress schedule storing means.

It is preferred that each ingress means includes means for storing a
25 plurality of transmission queues for transmission across the cross-bar. Advantageously, each ingress schedule storing means stores identities of said transmission queues, each transmission queue corresponding to a respective egress means identification number.

Preferably, each ingress means maintains a pointer into each ingress schedule storing means for identifying the transmission queue to be transmitted.

5 It is also preferred that each egress schedule storing means stores identities of ingress means addresses from which data is to be received. Advantageously, each egress means maintains a pointer into each egress schedule storing means for identifying an ingress means address from which data is to be received.

10 In accordance with another aspect of the present invention, there is provided a method of routing data using a switching arrangement as described above, the method comprising:-

- a) storing a plurality of transmission queue identities in each ingress schedule storing means;
- b) storing a plurality of ingress identities in each egress schedule
15 storing means;
- c) managing the contents of each ingress schedule storing means and each egress schedule storing means by providing ingress pointer means to reference one of said stored plurality of transmission queue identities and egress pointer means to reference one of said stored plurality of ingress
20 identities from which data is to be received; and
- d) at each cell transmit time, transmitting a cell from said referenced transmission queue in the ingress means and receiving the cell from said referenced ingress identity.

25 Preferably, step d) further comprises moving said ingress pointer and said egress pointer to the next location.

It is preferred that step c) comprises, for each ingress means, calculating cross-bar rates required to each egress means. Said cross-bar

rates may be calculated according to current traffic load and quality of service required.

Additionally, step c) further comprises calculating corresponding ingress and egress schedules which satisfy said calculated cross-bar rates.

- 5 The method further comprises the step of updating the ingress and egress schedule storing means with update messages relating to the calculated ingress and egress schedules.

The present invention provides a routing device having a plurality of ingress line function means, a plurality of egress line function means, a
10 backplane and a management card; each ingress line function means having: a schedule storing means, for storing a schedule of egress line function means addresses; a pointer storing means, for storing a pointer to an address held in the schedule storing means; and a queue storing means, for storing a plurality of ingress queues of cells for transmission across the
15 backplane, each one of the plurality of ingress queues corresponding uniquely to a predetermined one of the egress line function means; wherein the management card communicates configuration primitives to each of the plurality of ingress line function means and to each of the plurality of egress line function means, the configuration primitives providing updated
20 entries for the schedule.

The ingress line function means may be line interface cards. Likewise the egress line function means may also be line interface cards.

As a result of the present invention, simple schedules for cell transmission and reception across the cross-bar are sent to each LIC by the
25 controller. The controller is responsible for formulating the schedules to ensure that traffic and quality of service requirements are met. The LICs merely have to obey the schedules, a simple task. The principle benefits are reduced complexity of the LICs, and flexibility. New services can be

added to the router by software download without impact on the LIC hardware.

For a better understanding of the present invention, reference will now be made, by way of example only, to the accompanying drawings in
5 which:-

Figure 1 illustrates a terabit router architecture;

Figure 2 shows a cross-bar controller; and

Figure 3 shows a system for distributing schedules according to the present invention.

10 Although the present invention will be described with reference to ingress and egress forwarders, it will readily be understood that the present invention is not limited to the use of such forwarders. Any suitable line function means can be used for ingress and egress, for example, line interface cards (LICs).

15 Figure 1 illustrates a conventional terabit router architecture 100 in which packets arrive at ingress forwarders 102, 104, 106 via their input port(s) (not shown) and are routed across a cross-bar 110 to a correct egress forwarder 120 which transmits them across its output port(s) (not shown). Each ingress forwarder 102, 104, 106 maintains a separate packet queue for
20 each egress forwarder 120.

Ingress forwarder 102 has three queues q_{11} , q_{12} , q_{13} of data packets ready for transfer to three separate egress forwarders, only one of which is shown as 120, via the cross-bar 110. Similarly, three queues q_{21} , q_{22} , q_{23} and q_{31} , q_{32} , q_{33} are formed respectively in each of the ingress forwarders
25 104, 106. Although three queues are shown in each ingress forwarder 102, it will be appreciated that any number of queues can be present in each ingress forwarder 102, 104, 106, each queue corresponding to an egress means.

It will be appreciated that although only one egress forwarder 120 is shown in Figure 1, the number of egress forwarders will be the same as the number of ingress forwarders.

By way of explanation, a cell based cross-bar is characterised as follows:

a) Each ingress line function may be connected to any egress line functions.

b) Each ingress line function may only be connected to one egress line function at a time.

10 c) Each egress line function may only be connected to one ingress line function at a time.

d) All ingresses transmit in parallel across the cross-bar.

e) Data is transmitted across the cross-bar in small fixed sized cells, for example, a cell size is typically 64 octets.

15 f) Cell transmission is synchronised across all the ingress line functions. This means that for each cell cycle, each ingress line function starts transmitting the next cell at the same time.

g) The cross-bar is reconfigured at the end of every cell cycle.

As shown in Figure 1, packets of data arriving at the ingress forwarders 102, 104, 106 via their input port(s) (not shown) and are routed across the cross-bar 110 to the correct egress forwarders 120 which transmits them across its output port(s) (also not shown). Each ingress forwarder 102, 104, 106 maintains a separate packet queue for each egress forwarder 120, for example, q_{11} , q_{12} , q_{13} , q_{21} , q_{22} , q_{23} , q_{31} , q_{32} , q_{33} .

25 A conventional cell based cross-bar arrangement 200 is shown in Figure 2. The arrangement 200 comprises a plurality of ingress forwarders 210 and a plurality of egress forwarders 220 connected to a cross-bar or backplane 230. Here, each ingress forwarder 212, 214, 216, 218 may be

connected to one or more of the egress forwarders 222, 224, 226, 228.

However, as mentioned above, each ingress forwarder 212, 214, 216, 218 may only be connected to one egress forwarder 222, 224, 226, 228 at a time and each egress forwarder 222, 224, 226, 228 may only be connected to
5 one ingress forwarder at a time 212, 214, 216, 218.

The cross-bar arrangement 200 is controlled by a cross-bar controller 240 which is connected to each ingress forwarder 212, 214, 216, 218 via links 242, 244 and to each egress forwarder 222, 224, 226, 228 via link 246. The cross-bar controller 240 co-ordinates the transmission and
10 reception of cells via links 242, 244, 246.

Each ingress forwarder 212, 214, 216, 218 communicates traffic measurements and notifications for the use of the cross-bar controller 240. The cross-bar controller 240 allocates temporary connections between ingress forwarders 212, 214, 216, 218 and egress forwarders 222, 224, 226,
15 228 and informs the respective forwarders accordingly for each cell cycle in turn.

Figure 3 illustrates a system 300 in accordance with the present invention. The system 300 includes a plurality of ingress forwarders 310, a plurality of egress forwarders 320, a cross-bar 330, and a management card
20 340 which includes cross-bar controller means (not shown). The ingress forwarders 312, 314, 316 and egress forwarders 322, 324, 326 are shown as line interface cards (LICs), but may be implemented in other ways.

Although three ingress and three egress forwarders are shown, it will readily be appreciated that any suitable number of ingress and egress
25 forwarders may be utilised according to the particular application. In this arrangement, it is not essential that the number of ingress forwarders and egress forwarders are the same, for example, a multiplexer may have sixteen ingress forwarders and only four egress forwarders.

Each ingress forwarder 312, 314, 316 has a buffer (not shown in detail) containing a plurality of data queues $q_{11}, q_{12}, \dots, q_{1n}, q_{21}, q_{22}, \dots, q_{2n}$, and $q_{m1}, q_{m2}, \dots, q_{mn}$, respectively (where m is the number of ingress forwarders 310 and n is the number of data queues) feeding into respective
5 schedulers 352, 354, 356.

Each queue can be generally represented as q_{jk} where j indicates the ingress, k indicates the egress, and q_{jk} represents the packet queue at the ingress j for packets destined for egress k .

Each scheduler 352, 354, 356 is connected to a respective ingress
10 timetable 362, 364, 366 which controls the selection of data queue to be transferred to the cross-bar 330 in accordance with data from the management card 340.

Similarly, each egress forwarder 322, 324, 326 has a respective ingress selector 372, 374, 376 for receiving data from the cross-bar 330 and
15 for transferring it to a buffer (not shown in detail) containing a plurality of data queues $q'_{11}, q'_{12}, \dots, q'_{1p}, q'_{21}, q'_{22}, \dots, q'_{2p}$, and $q'_{r1}, q'_{r2}, \dots, q'_{rp}$, respectively (where r is the number of egress forwarders 320 and p is the number of data queues). Each ingress selector 372, 374, 376 is connected to a respective egress timetable 382, 384, 386 which provides information
20 as to which ingress forwarder to select in accordance with data from the management card 340. Each data queues $q'_{11}, q'_{12}, \dots, q'_{1p}, q'_{21}, q'_{22}, \dots, q'_{2p}$, and $q'_{r1}, q'_{r2}, \dots, q'_{rp}$ is connected to a further scheduler 392, 394, 396 which selects the appropriate queue to be output. This is, however, not essential to the present invention.

25 The operation of the method according to the present invention will now be described using the system 300 shown in Figure 3.

Each ingress LIC (ingress forwarder) 312, 314, 316 is associated with a respective schedule or timetable 362, 364, 366 governing the

transmission of cells by said ingress LIC. The schedule is in the form of a table whose entries are the identities of ingress LICs. Each ingress LIC maintains a pointer into the schedule. At each cell transmit time, the LIC transmits a cell from the queue identified in the entry referenced by the
5 pointer, and moves the pointer to the next location. The schedule is circular in the sense that when moving the pointer from the last entry, its next position is the first entry.

Each egress LIC (egress forwarder) 322, 324, 326 is also associated with a respective schedule or timetable 382, 384, 386 governing the
10 reception of cells by said egress LIC. The schedule is in the form of a table whose entries are the identities of ingress LICs from which to receive. An egress LIC maintains a pointer into the schedule. At each cell transmit time, the LIC will receive the cell from the ingress identified in the entry referenced by the pointer, and moves the pointer to the next location.
15 Again, the schedule is circular in the sense that when moving the pointer from the last entry, its next position is the first entry.

The management card 340 manages the contents of the ingress and egress schedules or timetables 362, 364, 366, 382, 384, 386. For each ingress LIC (ingress forwarder) 312, 314, 316, the management card 340
20 calculates the cross-bar rates required to each egress LIC (egress forwarder) 322, 324, 326. The cross-bar rates are calculated according to the current traffic load and required quality of service. The rates are calculated, for example, as described in the co-pending application mentioned above. Alternatively, the rates could be fixed.

25 Having calculated the rates, the management card 340 calculates corresponding ingress and egress schedules or timetables needed to satisfy them. It then updates the schedules or timetables using update messages known as configuration primitives on links 342, 344. Only modifications

to the schedules need be transmitted rather than the complete tables. This reduces the control traffic from the management card.

On receipt of the configuration primitives, the LICs or forwarders update the schedules as requested.

5 In order to avoid problems associated with updating the schedule, that is, reading and writing to the same point in schedule, the schedule may be partitioned so that reading and writing occur in different parts. Alternatively, two separate identical schedules may be provided – one schedule which is being read and another schedule which is being
10 configured.

This approach has the following advantages:-

First, complex rate calculation and schedule calculation algorithms can be implemented in software or on programmable devices on the management card 340, reducing development risk and cost.

15 Secondly, the hardware required to support the approach is very simple.

Thirdly, the communication channels between the LICs 310, 320 and the management card 340 do not require stringent delay performance since the primitives are not directly synchronised to the cell transmissions across
20 the backplane 330.

Fourthly, the communication channels 342, 344 between the LICs 310, 320 and the management card 340 require less bandwidth than would be the case if the management card were to reconfigure the cross-bar or backplane 330 every cell time.

25 Fifthly, the behaviour of the switch/router can be modified by changes to the software algorithms running on the management card 340. This has two consequences. Different applications can be supported by the same hardware. For example, the same hardware could support both

internet protocol (IP) routing and circuit switching, including asynchronous transfer mode (ATM) and the related synchronous transfer mode (STM). Furthermore, optimisations and refinements to the algorithms can easily be implemented.

- 5 It will be readily understood that although the preceding discussion has been in terms of terabit routers, the apparatus of the present invention are capable of implementation in a wide variety of switching devices, including switches and routers, and that these devices can be either purely electronic, part electronic/part optical or optical in nature.

CLAIMS:

1. A switching arrangement having:-
a cross-bar;
a plurality of ingress means connected to an input side of the cross-bar, each ingress means including an ingress schedule storing means;
a plurality of egress means connected to an output side of the cross-bar, each egress means including an egress schedule storing means; and
a management card which communicates configuration primitives to each of the plurality of ingress means and to each of the plurality of egress means, the configuration primitives providing updated entries for ingress and egress schedule storing means.
2. An arrangement according to claim 1, wherein each ingress means includes means for storing a plurality of transmission queues for transmission across the cross-bar.
3. An arrangement according to claim 2, wherein each ingress schedule storing means stores identities of said transmission queues, each transmission queue corresponding to a respective egress means identification number.
4. An arrangement according to claim 3, wherein each ingress means maintains a pointer into each ingress schedule storing means for identifying the transmission queue to be transmitted.

5. An arrangement according to any one of the preceding claims, wherein each egress schedule storing means stores identities of ingress means addresses from which data is to be received.
6. An arrangement according to claim 5, wherein each egress means maintains a pointer into each egress schedule storing means for identifying an ingress means address from which data is to be received.
7. A method of routing data using a switching arrangement according to any one of the preceding claims, the method comprising:-
 - a) storing a plurality of transmission queue identities in each ingress schedule storing means;
 - b) storing a plurality of ingress identities in each egress schedule storing means;
 - c) managing the contents of each ingress schedule storing means and each egress schedule storing means by providing ingress pointer means to reference one of said stored plurality of transmission queue identities and egress pointer means to reference one of said stored plurality of ingress identities from which data is to be received; and
 - d) at each cell transmit time, transmitting a cell from said referenced transmission queue in the ingress means and receiving the cell from said referenced ingress identity.
8. A method according to claim 7, wherein step d) further comprises moving said ingress pointer and said egress pointer to the next location.

9. A method according to claim 8, wherein step d) further comprises, when each pointer has reached the last location, moving to the first location.
10. A method according to any one of claims 7 to 9, wherein step c) comprises, for each ingress means, calculating cross-bar rates required to each egress means.
11. A method according to claim 10, wherein said cross-bar rates are calculated according to current traffic load and quality of service required.
12. A method according to claim 10 or 11, wherein step c) further comprises calculating corresponding ingress and egress schedules which satisfy said calculated cross-bar rates.
13. A method according to claim 12, further comprising the step of updating the ingress and egress schedule storing means with update messages relating to the calculated ingress and egress schedules.
14. A switching arrangement substantially as hereinbefore described with reference to Figure 3 of the accompanying drawings.
15. A method of routing data substantially as hereinbefore described with reference to Figure 3 of the accompanying drawings.

ABSTRACT

IMPROVEMENTS IN OR RELATING TO SWITCHING DEVICES

The invention relates to a switching device (300) having a plurality of line interface cards (LICs) (310), a plurality of egress LICs (320), a cross-bar (330) and a management card (340). The switching device (300) routes fixed sized cells of data across the cross-bar (330). The cells comprise fragments of variable length data packets. Each ingress LIC (312, 314, 316) is associated with a respective schedule or timetable (362, 364, 366) governing the transmission of cells by the ingress LIC. Similarly, each egress LIC (322, 324, 326) is also associated with a respective schedule or timetable (382, 384, 386) governing the reception of cells by said egress LIC. The schedules are in the form of a table whose entries are the identities of transmission queues corresponding to a respective egress LIC identification number (for ingress LICs) and ingress LICs from which to receive (egress LICs). Each ingress and egress LIC maintains a pointer into its associated schedule. At each cell transmit time, each ingress LIC transmits a cell from the queue identified by the pointer and each egress LIC will receive the cell from the ingress LIC identified in the entry referenced by the pointer. The pointers then move to the next location. Each schedule is circular in the sense that when moving the pointer from the last entry, its next position is the first entry.

(Fig. 3)

THIS PAGE BLANK (USPTO)

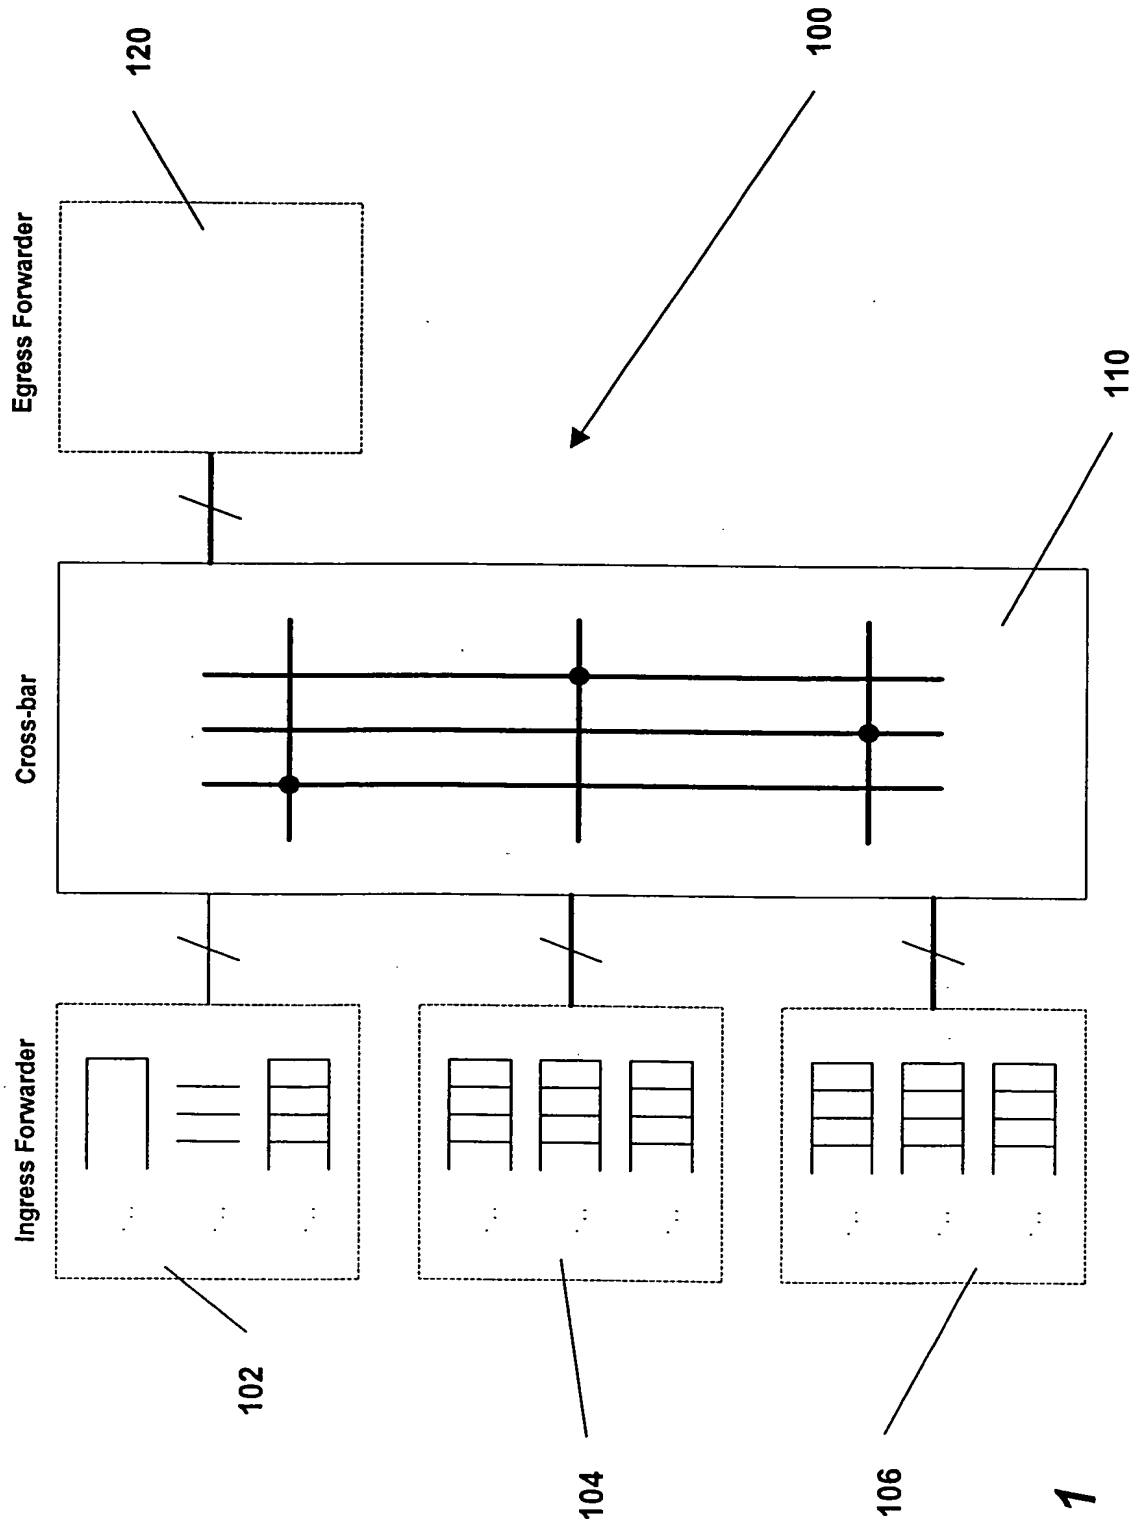



Fig. 1

THIS PAGE BLANK (USPTO)

**Fig. 2**

THIS PAGE BLANK (USPTO)




Fig. 3

THIS PAGE BLANK (USPTO)